

Method and Apparatus for Using Thermal Infrared for Face Recognition

Field of the Invention

[001] This application claims the benefit of U.S. Provisional Application No. 60/393,118, filed July 3, 2002, the disclosure of which is hereby incorporated herein by reference.

[002] The invention described herein relates to the use of thermal infrared imaging for face recognition.

Background of the Invention

[003] Over the past several years the development of face recognition systems has been receiving increased attention as having the potential for providing a non-invasive way of improving security systems and homeland defense. Such systems may be used for applications such as access control to restricted areas, by either control of physical entry into a building, room, vault or an outdoor area or electronically such as to a computer system or ATM. Another application of such systems is identification of individuals on a known watchlist, which can consist of but is not limited to, known criminals, terrorists, or casino cardcounters. For identification a face recognition system produces a rank ordering of known individuals that closely match an unknown subject. If there is an identification matched ranking of N (e.g., $N=10$) or less with a known malevolent individual, then the unknown subject can either be detained or taken to a secondary procedure where further information is solicited. Another set of applications include

surveillance and monitoring of scenes whereby the identity of individuals present in a scene is periodically verified.

[004] Existing end-to-end systems that detect and recognize faces of individuals at a distance are exclusively performed with visible light video cameras. The influence of varying ambient illumination on systems using visible imagery is well-known to be one of the major limiting factors for recognition performance [Wilder, Joseph and Phillips, P. Jonathon and Jiang, Cunhong and Wiener, Stephen, "Comparison of Visible and Infra-Red Imagery for Face Recognition," Proceedings of 2nd International Conference on Automatic Face & Gesture Recognition, pp. 182--187, Killington, VT, 1996; Adini, Yael and Moses, Yael and Ullman, Shimon, "Face Recognition: The Problem of Compensating for Changes in Illumination Direction," IEEE Transactions on Pattern Analysis and Machine Intelligence, Volume 19, No. 7, pp. 721--732, July, 1997]. A variety of methods compensating for variations in illumination have been studied in order to boost recognition performance, including histogram equalization, Laplacian transforms, Gabor transforms, logarithmic transforms, and 3-D shape-based methods. These techniques aim at reducing the within-class variability introduced by changes in illumination, which has been shown to be often larger than the between-class variability in the data, thus severely affecting classification performance. System performance, particularly outdoors where illumination is dynamic, is problematic with existing systems.

[005] Face recognition in the thermal infrared domain has received relatively little attention compared with recognition systems using visible-spectrum imagery. Original tentative analyses have focused mostly on validating the thermal imagery of faces as a

valid biometric [Prokoski, F. J., “History, Current Status, and Future of Infrared Identification, Proceedings IEEE Workshop on Computer Vision Beyond the Visible Spectrum: Methods and Applications, Hilton Head, June 2000; Wilder, Joseph and Phillips, P. Jonathon and Jiang, Cunhong and Wiener, Stephen, “Comparison of Visible and Infra-Red Imagery for Face Recognition,” Proceedings of 2nd International Conference on Automatic Face & Gesture Recognition, pp. 182--187, Killington, VT, 1996]. The lower interest level in infrared imagery has been based in part on the following factors: much higher cost of thermal sensors versus visible video equipment, lower image resolution, higher image noise, and lack of widely available data sets. These historical objections are becoming less relevant as infrared imaging technology advances, making it attractive to consider thermal sensors in the context of face recognition.

Summary of the Invention

[006] Thermal infrared imagery of faces is nearly invariant to changes in ambient illumination [Wolff, L. and Socolinsky, D. and Eveland, C., “Quantitative Measurement of Illumination Invariance for Face Recognition Using Thermal Infrared Imagery, Proceedings CVBVS, Kauai, Dec. 2001]. Consequently, no compensation is necessary, and within-class variability is significantly lower than that observed in visible imagery [Wolff, L. and Socolinsky, D. and Eveland, C., Quantitative Measurement of Illumination Invariance for Face Recognition Using Thermal Infrared Imagery,” Proceedings CVBVS, Kauai, Dec. 2001]. It is well-known that for visible video the set of images of a given face acquired under all possible illumination conditions is a subspace of the vector space of images of fixed dimensions. In sharp contrast to this, the set of thermal infrared images of a face under all possible imaging conditions is contained in a bounded set. It follows

that under general conditions lower within-class variation can be expected for thermal infrared images of faces than their visible counterpart. It remains to be demonstrated that there is sufficient between-class variability to ensure high discrimination, but the combined use of both reflective spectrum and thermal infrared imagery provides even more accurate discrimination [Socolinsky, D., Wolff, L., Neuheisel, J., and Eveland, C. "Illumination Invariant Face Recognition Using Thermal Infrared Imagery," Computer Vision and Pattern Recognition, Kauai, December 2001 & Socolinsky, D. and Selinger A., "A Comparative Analysis of Face Recognition Performance with Visible and Thermal Infrared Imagery," ICPR '02, Quebec, August 2002].

[007] A key aspect of a face recognition system is to be able to store and match representation templates of faces. Creating face representation templates from both reflective spectrum and thermal infrared imagery provides significant advantages. In the present invention both video imagery sensed in a sub-spectrum of the reflective domain and video imagery sensed in a sub-spectrum of the thermal infrared domain are used to detect a face, create a face representation template, and match/compare the face representation template of an unknown individual with a stored database or gallery of face templates. This can be applied to a variety of uses for face recognition systems. Reflective domain and thermal infrared domain imagery have a low degree of correlation with respect to the phenomenological information that they sense from a scene. This makes such imaging modalities highly complementary in the additional information they provide each other, which is particularly useful for face recognition systems.

[008] This invention also includes sensor technology that acquires both reflective spectrum (e.g., visible) imagery and thermal infrared imagery. This can either consist of

separate visible and thermal infrared sensors or integrated visible/thermal infrared sensors. Cameras that use CCD, CMOS and CID focal plane arrays (FPA) have sensitivity typically in the 0.4-1.0 micron wavelength range spanning the visible and near-infrared spectrums. The InGaAs FPA made by Sensors Unlimited has a sensitivity range typically 0.9-1.7microns. Cooled InSb FPA has good sensitivity in the 3-5micron thermal infrared range, while cooled MCT or QWIP FPA, and uncooled microbolometer FPA, have good sensitivity in the 8-12 micron range. It is important to apply proper non-uniformity correction (NUC) and radiometric calibration procedures to thermal infrared imagery.

Brief Description of the Drawings

[009] FIG. 1 is a taxonomy of important regions of the electromagnetic spectrum, defining the Reflective and Thermal Infrared domains (spectrums) discussed herein;

[0010] FIG. 2 is a block diagram illustrating thermal infrared video imaging in conjunction with reflected video imaging for creating a face recognition representation template for use in an end-to-end system;

[0011] FIG. 3 (a) illustrates a method for non-uniformity correction (NUC) of thermal infrared imagery using a constant temperature flag; (b) illustrates NUC or radiometric calibration of thermal infrared imagery using a blackbody source; and FIG. 3 (c) illustrates an example of how pattern noise is removed from a thermal infrared image;

[0012] FIG. 4 is a flow chart of a face detection monitoring loop

[0013] FIG. 5 (a) is a flow chart of a face representation template for access control; and FIG. 5 (b) is a flow chart of a face representation template for verifying identity of individuals being periodically monitored;

[0014] FIG. 6 is a flow chart of a face representation template for identification;

[0015] FIG. 7 is a block diagram showing creation of a face representation template from masked subregions of images from the reflective spectrum and from the thermal infrared spectrum;

[0016] FIG. 8 is a block diagram showing hardware apparatus required to implement a system for creating face representation template from reflective and thermal infrared imagery;

[0017] FIGS. 9 (a) through 9 (e) illustrate a number of configurations of reflective spectrum and thermal infrared spectrum imaging sensors at various levels of integration; and

[0018] FIG. 10 is a block diagram of apparatus for a face recognition system with multiple sets of video camera(s) for monitoring various respective locations at once.

Description of the Invention and Preferred Embodiments

[0019] It should be noted that the term *infrared* as used in the literature does not always refer to thermal infrared, and in fact as shown in Figure 1 by the spectrum **90** there are important imaging sub-spectrums of the infrared spectrum that primarily record reflective phenomenology, such as the near-infrared **92** and the shortwave infrared **94** (SWIR). Image fusion described herein refers to fusion of images taken from a sub-spectrum of the reflective domain **101**, and a sub-spectrum of the thermally emissive (i.e., thermal infrared) domain **102** as specified in Figure 1.

[0020] Figure 2 illustrates a face recognition system **200** receiving two video imaging streams **210** and **220** from suitable sensors or cameras **250** and **260**, respectively. One of the video imaging streams **210** is produced from sensing in the reflective domain **101**,

while the other video imaging stream **220** is produced from sensing in the thermal infrared domain **102**. In a preferred embodiment, the reflective video stream **210** is in the visible spectrum **103** and is produced from a CCD focal plane array (FPA) in camera **250**, while the thermal infrared video stream **220** is in the LWIR spectrum **104** and is produced from a microbolometer focal plane array (FPA) in camera **260**. Although two cameras are illustrated, it will be understood that the respective FPAs may be in a single camera. In a preferred embodiment, image pixels produced from the reflective video imaging stream are spatially co-registered with image pixels produced from the thermal infrared video stream. Typically there is more spatial resolution available from sensing in the reflective domain than there is available from sensing in the thermal infrared domain. In a preferred embodiment, multiple pixels from the reflective domain are assigned/co-registered to a single pixel in the thermal infrared domain.

[0021] Imagery produced from most thermal infrared focal plane arrays, as illustrated at **310** in Fig. 3 (c), experience a significant amount of pattern noise with variable gain and offset for different pixels across the focal plane. Non-uniformity correction (NUC) image pre-processing, illustrated at **221** in Fig. 2, is used to clarify the image as illustrated at **204**, and is required for thermal infrared imagery prior to use for face detection and further creation of a face representation template. Figures 3(a) and (b) show apparatus that can be used to perform this. Figure 3(a) shows a common apparatus for performing a one-point NUC using a thermally opaque flag **301**, usually made of metal which is at a constant temperature throughout, that periodically slides in front of the thermal infrared focal plane array (FPA) **300**, which may be in camera **260**, actuated by a servo or solenoid mechanism **302**. Image processing insures that image gray values

outputted at **304** or **306** (Fig. 2) by the thermal infrared FPA are exactly equal for all pixels across the focal plane when receiving equivalent thermal emission from a scene element as from the flag when calibration took place. This is also called a *one-point* NUC. The gray value response of a thermal infrared camera is linear with respect to the amount of incident thermal radiation. The slope of the responsivity line is called the *gain* and the y-intercept is the *offset*. As mentioned, gain and offset for each pixel on a thermal infrared FPA vary significantly across the array. A one-point NUC insures that all responsivity lines respective to all pixels intersect at the thermal emission value for the flag.

[0022] A one-point NUC can also be achieved by using a blackbody source (e.g., a Mikron model 350), illustrated at **303** in Fig. 3 (b), which not only has exactly uniform thermal emission spatially across a flat black surface, but the temperature of this flat black surface can be accurately controlled. Two separate thermal infrared images taken respectively for blackbody temperatures T_1 and T_2 produces a *two-point* NUC which establishes the direct linear relationship between the gray value response at a pixel and the absolute amount of thermal emission from the corresponding scene element. That is the gain and offset are precisely known for each pixel. A NUC such as element **221** in Fig. 2, can either refer to one-point or two-point procedure.

[0023] In the reflective domain, for InGaAs FPAs sensitive to SWIR radiation there is also a necessity to perform at least a one-point NUC **211** producing outputs **322** and **324**. Other reflective domain FPAs, particularly scientific grade FPAs, may also require a one-point NUC. Typically, this is performed using the flag apparatus **301** shown in Figure 3(a), such as for the Indigo Merlin InGaAs camera. However, for most CCD and CMOS

FPAs, NUC is not an issue and **211** can be bypassed, as illustrated by lead lines **330** and **332**.

[0024] While a two-point NUC in the thermal infrared provides non-uniformity correction, the relationship back to a physical property of the imaged object---its emissivity and temperature---provides the further advantage of data where environmental factors contribute to a much lesser degree to within-class variability. An added bonus of using a two-point NUC for thermal infrared is that it simplifies the problem of skin detection in cluttered scenes [Eveland, C., Socolinsky, D., and Wolff, L. "Tracking Human Faces in Infrared Video," CVPR Workshop on Computer Vision Beyond the Visible Spectrum, Kauai, December 2001]. The range of human body temperature is quite small, varying from 36 deg. C to 38 deg. C. We have found that skin temperature at 22 deg. C ambient room temperature to also have a small variable range from about 26 deg. C to 29 deg. C. Two-point NUC makes it possible to perform an initial segmentation of skin pixels in the correct temperature range.

[0025] One can achieve marginally higher precision by taking blackbody measurements at multiple temperatures and obtaining the gains and offsets by least squares regression. For the case of thermal images of human faces, each of the two fixed temperatures are below and above skin temperature, respectively, to obtain the highest quality calibration for skin levels of thermal emission.

[0026] It should be noted that calibration has a limited life span. If a NUC is performed on a thermal infrared camera indoors, taking it outdoors where there is a significant ambient temperature difference will cause the offsets of individual pixels to change. Therefore, a NUC must be performed again. This effect is due mostly to temperature

variations of the lens optics and FPA. Also, if two separate data collections (i.e., a set of image acquisitions of individuals) are performed with different thermal infrared cameras, even with the exact same model number, identical camera settings and under the exact same environmental conditions, the gain and offset of corresponding pixels between these sensors will differ since no two thermal infrared focal plane arrays are ever identical. Yet another example: if two data collections are performed one year apart, with the same thermal infrared camera, it is very likely that gain and offset characteristics will have changed. Two-point NUC standardizes all thermal infrared data collections, whether they are taken under different environmental conditions or with different cameras or at different times. Since pixelwise gray values for a thermal infrared image are directly related to the thermal emission power of the imaged scene, this provides a standardized thermal IR biometric signature for humans. The most beneficial images for face recognition algorithms are not arrays of gray values, but rather of corresponding thermal emission values. This is one critical difference between thermal and visible imaging for face recognition: the inability to relate visible intensities to intrinsic properties of the object makes it impossible to use absolute gray values as a reliable recognition feature.

[0027] In order to initiate the face recognition system of the present invention for any application, a face must be detected in the scene. In accordance with the method of the present invention (Fig. 4), a scene is continuously monitored, at **401**, until a face is detected, at **402**. When a face is detected by detector **230** (Fig. 2), eyes for the face are then detected at detector **212** and this detector output is used for geometric normalization of the reflective video image as indicated at **214**. Similarly, the thermal infrared video image is geometrically normalized at **224**. For geometric normalization at **214** and **224**,

the locations of the eyes are used to affinely map the images to a standard geometry, with fixed eye locations. These images are then sub-sampled and subsequently cropped with corresponding masks **215** and **225** that remove all but the inner face, thus eliminating the effect of background, hair, etc. After masking image normalization is performed at **216** and **226**, that statistically demeans and Gaussian normalizes the gray values of the masked image.

[0028] Many face recognition (and general pattern classification) algorithms can be divided into two stages: first, a feature representation such as that carried out at **217**, **227**, and **232** in Fig. 2 and resulting from an image input, or probe, and second, a similarity computation illustrated at **240** in Fig. 2 and carried out using the steps **500**, **550**, **600** illustrated in Figs. 5 (a), 5 (b), and 6 to match the probe image with a gallery of images in a database. The feature representation stage is responsible for encoding the stimulus to be classified in the form of a sequence of numbers, or a feature vector. This vector is usually of fixed length, and the mapping process from stimuli to feature vectors is fixed when a system is initiated. Once both probes (unknown faces) and gallery (known faces) data have been mapped to the feature vector space, classification proceeds by considering the similarity between the probe and each gallery exemplar. This similarity can be computed with respect to multiple measures, yielding different performance characteristics. It is often the case that a feature representation is constructed to be optimal for a specific similarity measure. However, it is possible that a different similarity measure yields better classification performance when paired with that feature representation.

[0029] For normalization elements **216** and **226**, preprocessed face images respectively in the reflective spectrum and in the thermal infrared spectrum, in the form of n -tuples, are mapped into a k -dimensional feature space (where $1 \leq k < \infty$) via a linear, affine or fully non-linear map. These k -tuples are known as *templates*. The preferred embodiment of step **240** incorporates a nearest-neighbor classification with respect to a norm on R^k . Alternative embodiments can incorporate neural networks or support vector machines (SVM) or other classifiers. There are various existing methods for generating templates; for example Turk, M. and Pentland, A., "Eigenfaces for Recognition," J. Cognitive Neuroscience, Volume 3, Pages 71--86, 1991; Belhumeur, P. and Hespanha, J. and Kriegman, D., "Eigenfaces vs. Fisherfaces: Recognition Using Class Specific Linear Projection," IEEE Transactions PAMI, Volume 19, No. 7, Pages 711--720, July, 1997; Penev, P. and Attick, J., "Local Feature Analysis: A general statistical theory for object representation, Network: Computation in Neural Systems," Volume 7, No. 3, Pages 477--500, 1996; R. J. Michaels and T. Boulton, "Efficient evaluation of classification and recognition systems," Proceedings of IEEE Computer Vision and Pattern Recognition, Kauai, HI, Dec. 2001; A. J. Bell and T. J. Sejnowski, "An Information-Maximization Approach to Blind Separation and Blind Deconvolution," Neural Computation Volume 7, Number 6, Pages 1129--1159, 1995; C. Liu and H. Wechsler, "Comparative Assessment of Independent Component Analysis (ICA) for Face Recognition," Proceedings of the Second Int. Conf. on Audio- and Video-based Biometric Person Authentication, Washington, DC, March 1999; P. Comon, "Independent component analysis: a new concept?," Signal Processing, Volume 36, Number 3, Pages 287--314, 1994, as well as others.

[0030] When step 240 is a nearest neighbor classifier, there are multiple choices as to the distance function used for classification (or dissimilarity). A *dissimilarity measure* is a function $\delta : R^k \times R^k \rightarrow R$ such that $\arg \min \delta(v, w) = v$, for all $v \in R^k$. Although this is not strictly necessary, it is customary to assume that δ is continuous in each variable. Multiple dissimilarity measures may be used.

[0031] A combination of classifiers can often perform better than any one of its individual component classifiers. In fact, there is a rich literature on the combination of classifiers for identity verification, mostly geared towards combining voice and fingerprint, or voice and face biometrics (e.g. J. Big'un and B. Duc and F. Smeraldi and S. Fischer and A. Makarov, "Multi-Modal Person Authentication, Proceedings of Face Recognition: From Theory to Applications," Stirling, UK, NATO Advanced Study Institute, Jul, 1997; B. Achermann and H. Bunke, "Combination of Classifiers on the Decision Level for Face Recognition," Institute of Computer Science and Applied Mathematics, University of Bern, Number IAM-96-002, Bern, Switzerland, Jan, 1996). The degree to which combining the results of two or more classifiers improves performance is highly dependent on the degree of correlation among classifier decisions. Combining several highly correlated classifiers normally has no effect beyond that of increasing system complexity, whereas fusing experts with low correlation can dramatically improve performance. Some results and further considerations on fusing face recognition algorithms on visible imagery can be found in W. S. Yambor and B. A. Draper and J. R. Beveridge, "Analyzing PCA-based Face Recognition Algorithms: Eigenvector Selection and Distance Measures," Proceeding 2nd Workshop on Empirical Evaluation in Computer Vision, Dublin, Ireland, 2000. In the case of coregistered

visible/thermal imagery, a perfect opportunity for classifier fusion is available, since from physical principles it is known that there is very low correlation between the data in the two modalities:

[0032] In a preferred embodiment of step **240** a weighted combination of normalized scores from each classifier corresponding to either full face or sub-windows from one or more of the reflective/thermal infrared imaging modalities is used. These scores could be distance or dissimilarity values in the case of nearest-neighbor classifiers, neuron activation strength in the case of neural networks, or distances to a separating hyperplane when used with SVM. The weights can be adaptably varied to account for relative classifier performance, image degradation, occlusion, expression variation, etc. Individual scores may be normalized to account for variability in statistical properties of scores output by different classifiers.

[0033] As one example, a simple adaptive weighting scheme is introduced that yields a single distance-like score from the dissimilarity score returned by two classifiers, each one acting on one modality from a visible/thermal pair. Let G^v and G^l be the visible and thermal infrared image galleries, respectively, and p^v and p^l be the visible and thermal infrared components of a bi-modal probe image. Lastly, let δ^v and δ^l be the dissimilarity measures on the respective feature vector spaces corresponding to the recognition algorithms used on each modality. For any feature vector $(g^v, g^l) \in G^v \times G^l$, a combined dissimilarity score is defined by:

$$\delta^f((p^v, p^l), (g^v, g^l)) = \frac{1}{3} \left[\frac{\delta^v(p^v, g^v)}{\varpi^v} + 2 \frac{\delta^l(p^l, g^l)}{\varpi^l} \right] \quad (\text{Eq. 1})$$

where ϖ^v and ϖ^l are the median values of the sets $\{\delta^v(p^v, G^v)\}$ and $\{\delta^l(p^l, G^l)\}$, respectively. The median factors dividing each term insure that dissimilarity scores on widely different scales can be combined without one overwhelming the other. A fused face score from reflective spectrum and thermal infrared spectrum imagery at 232 can be accordingly made, with associated classifiers to perform matching decisions 240, 500, 550 and 600. This is only an example and not a preferred embodiment.

[0034] Figure 7 shows that fusion between reflective spectrum and thermal infrared imagery can go beyond just using the entire face image for each modality. Martinez, A.M., "Representing Imprecisely Localized, Partially Occluded, and Expression Invariant Faces from a Single Sample per Class," IEEE Transactions on PAMI, Vol. 24, No. 6, June 2002 discusses a methodology by which multiple sub-windows/sub-regions of face images can each be reduced to a representation (i.e., template), and then these representations can be combined. This is particularly useful in the presence of temporary occlusions and variant face expressions. This methodology is extended, by the present invention, to fusion of reflective spectrum and thermal infrared imagery. When performing this method, steps 715, 716 and 717 correspond to steps 215, 216 and 217, respectively; and steps 725, 726 and 727 correspond to steps 225, 226 and 227, respectively. Each one of the face feature templates in 717 and 727 corresponds to a selected sub-window image region. The sub-windowing scheme for the reflective spectrum image need not be the same as the sub-windowing scheme for the thermal infrared image.

[0035] Hardware apparatus for implementing the methodology taught here as illustrated diagrammatically in Fig. 8, consists of a camera configuration 801 (or simply, camera)

capable of simultaneously acquiring a reflective spectrum image and a thermal infrared spectrum image and able to output this imagery at **802** either in an analog or digital mode. Figure 9 illustrates a number of specific camera configurations that can be used for determining camera **801**. The camera **801** interfaces to a computer system **820** by connecting with an interface card **822** which could be a digital interface card if the camera output is digital or an analog/digital converter if the camera output is analog. Interface card **822** enables reflective spectrum and thermal infrared images to be placed in computer memory **824**. This can be stored in disk memory **826**. Creation of the face representation template and subsequent use of this template to compare and match with existing templates for face recognition applications is performed by software operating on computer system **820**. ACE face recognition software, version 1.0 by Equinox, Inc., 9 West 57th Street, NY, NY, is one example of this. This software operating on computer **820** also performs the face detection process monitoring when a face is present in the scene of interest, and determines when to acquire imagery of unknown subjects, store imagery, create face representation templates, and perform matching comparisons for specified face recognition applications.

[0036] Figure 9 illustrates a number of ways to configure a camera sensor that simultaneously acquires reflective spectrum and thermal infrared imagery. At the least, an FPA **910** that senses in the reflective spectrum (e.g., CCD 0.4-1.0micron, InGaAs 0.9-1.7 micron), and an FPA **920** that senses in the thermal infrared (e.g., InSb 3.0-5.0 microns, MCT 8-12microns, microbolometer 8-14microns) are required. Figure 9(a) shows two separate cameras with respectively separate electronics and separate optics. These two cameras are respectively viewing reflective domain radiation **911** and thermal

infrared domain radiation **921** from the same elements of a scene, although from possibly different viewing directions. Figure 9(b) shows a boresighted configuration **930** for the same respective cameras **910** and **920**. In this configuration the optical axis for each camera lens is parallel, and viewing differences between the two cameras only exists with respect to the translational baseline segment determined by the optical centers for each camera. This produces some translational disparity between the respective image pixels for the reflective domain and thermal infrared domain FPAs.

[0037] Figure 9(c) shows configurations **932** and **934** using the same two separate cameras **910** and **920** as in 9(a,b) but incorporating a dichroic beamsplitter **940** that takes radiation **936** from a scene and either transmits thermal infrared domain radiation and reflects reflected domain radiation, or vice versa. A dichroic beamsplitter used in this fashion further reduces the baseline displacement between the reflective domain and thermal infrared domain cameras.

[0038] Figure 9(d) shows an integrated camera **944** having two FPAs **946** and **948**, respectively, sensing in the reflective domain and in the thermal infrared domain. An important difference between 9(c) and 9(d) is that it includes a dichroic beamsplitter **942** completely behind all common focusing optics **950**. This completely eliminates depth dependent disparity between the reflective domain and thermal infrared domain FPAs. Figure 9(e) depicts a camera **958** with a hybrid FPA **960** capable of sensing both a subspectrum of reflective domain and thermal infrared radiation **964** with a focusing lens **962**.

[0039] Figure 10 shows hardware apparatus for a distributed system of cameras supporting face recognition applications where unknown subjects are being monitored at

multiple remote locations, such as in a large building complex or even in a home with multiple rooms. Each reflective spectrum/thermal infrared imaging camera **1001** at a separate remote location is connected via a corresponding computer interface card **1002** to its own dedicated computer board **1003** such as a PC-104+ with PentiumIII. The task of detecting a face at each respective remote location, acquiring the reflective spectrum/thermal infrared imagery of the detected face and reducing to a face representation template is done by the dedicated computer at that remote location. The face representation template is then sent over a high-speed line **1010** (e.g., ethernet) to a main computer **1020** for comparison matching and implementation of the particular face recognition application required.

[0040] Although the invention has been described in terms of various preferred embodiments, it will be understood that numerous variations and modifications may be made without departing from the true spirit and scope thereof, as set forth in the following claims.